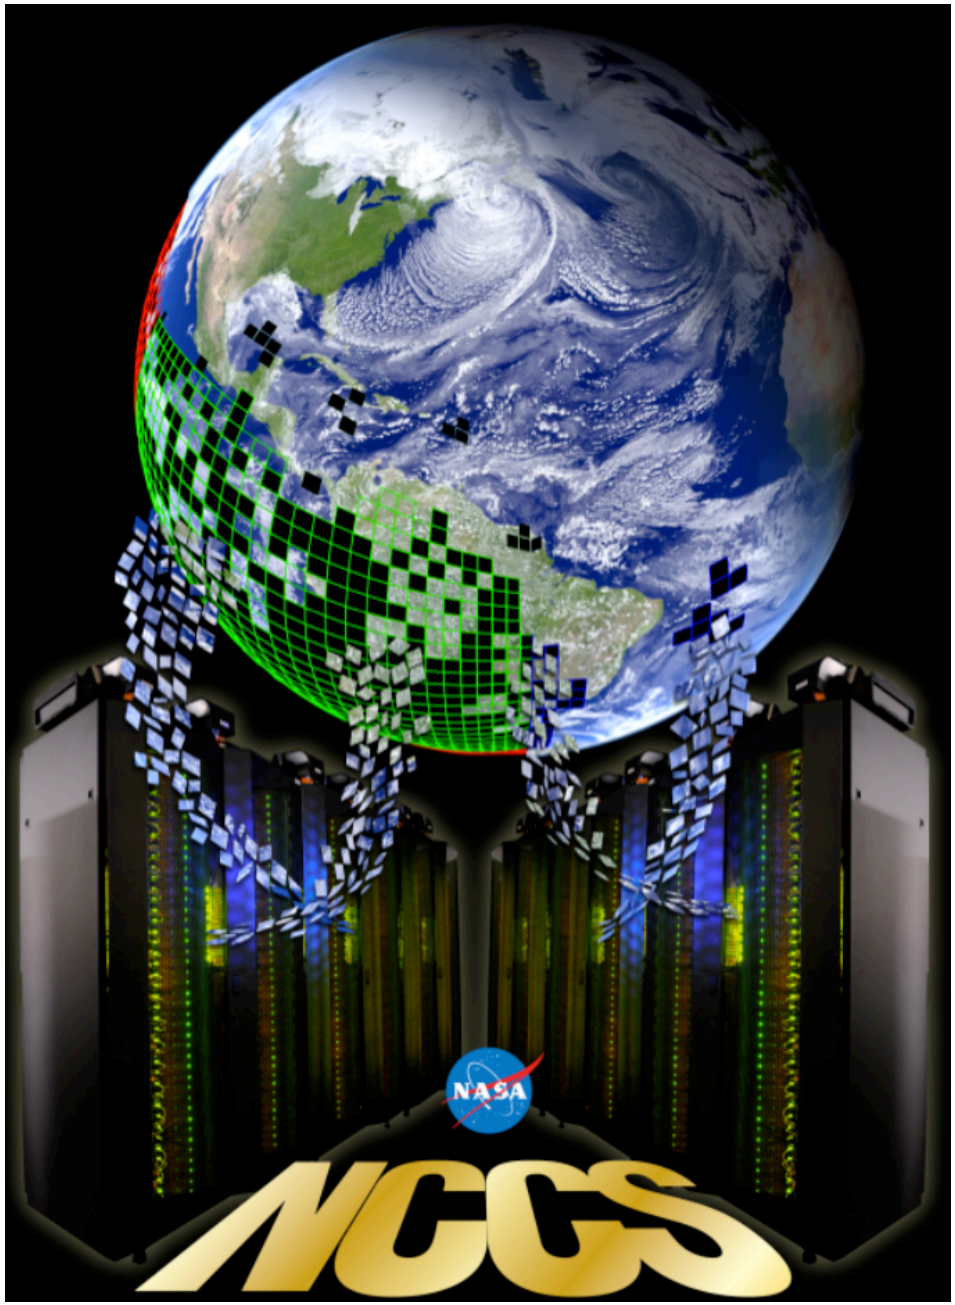

Climate Simulation at Goddard

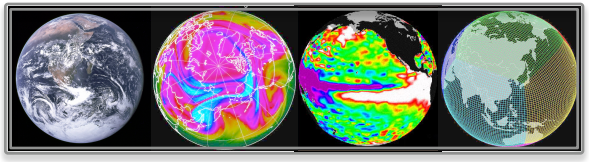

NASA High-end Computing Support for AR5

Phil Webster & Harper Pryor
Computational and Information Science & Technology Office
NASA Center for Computational Science (NCCS)

Goddard Space Flight Center

GMAO Modelling for AR5

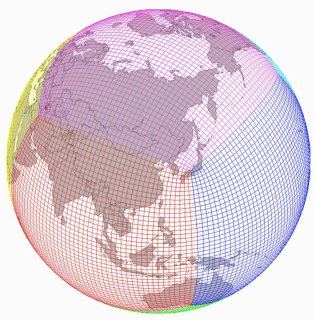

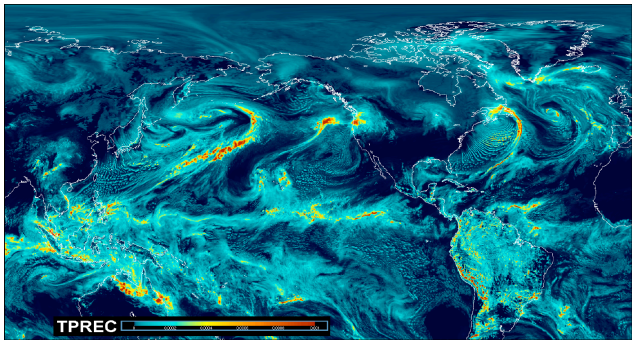

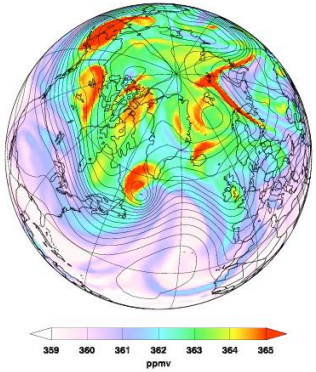

- Decadal prediction with GEOS-5 AOGCM
  - Coupled A-O-L initialization
  - **10-years:** 1º agcm; 1/2º ogcm; 10 mem, 10 cases (every 5 years, 1960-2005)
  - **30-years:** 1º agcm; 1/2º ogcm; 10 mem, 3 cases (1960, 1980, 2005)

- Atmospheric chemistry & ozone changes (GEOS-CCMv2)
  - Time-slice runs focused on 2030-2040 (stratospheric chemistry)
  - Historical runs, including uncertainty

- Atmospheric chemistry (GEOS-CCMv3 with AOGCM)
  - low resolution with full chemistry

~ 2200 simulation years

GISS Modelling for AR5

Gavin Schmidt, GISS

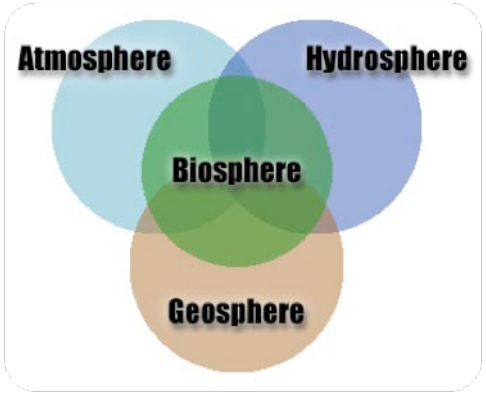

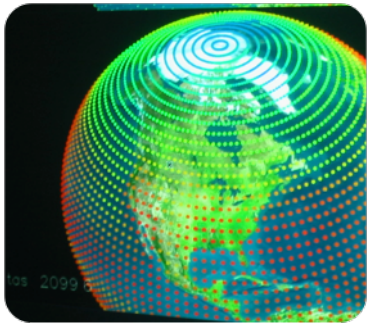

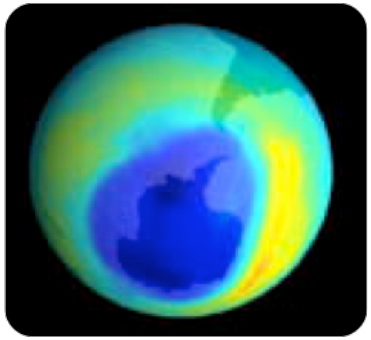

ModelE AOGCM Current status/configuration

- Model physics frozen:
  - Improved clouds, sea ice, new dynamical core options (includes Fvcore – collaboration with GSFC)
- Atmosphere: 2x2.5, 40 layers. Hi-resolution version using Cubed-Sphere C90 (~1x1)
- Oceans: Two versions (spin-ups underway):
  - HYCOM – tri-polar grid (1x1 – equatorial refinement to 0.2 lat), hybrid isopycnic
  - Russell – lat-lon (~1x1.25), z*-levels

ESM current status/configuration

- Fully interactive aerosols, stratosphere/troposphere gas phase chemistry (Koch, Menon, Shindell)
- Indirect effects (AIE2 + BC/albedo) included
- Carbon cycle includes NOBM (ocean) CC, Ent (terrestrial) CC

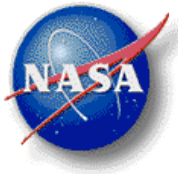

NASA High Performance Computing

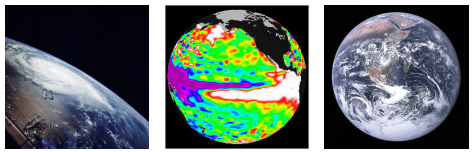

NCCS

- **NCCS at Goddard Space Flight Center**
  - Focused on Climate and Weather Research in the Earth Science Division of the Science Mission Directorate
    - Support code development
    - Environment for running models in production mode
    - Capacity computing for large, complex models
    - Analysis & visualization environments
- **NAS at Ames Research Center**
  - Supports all Mission Directorates
    - For Earth Science: Capability runs for test & validation of next generation models

NCCS Data Centric Climate Simulation Environment

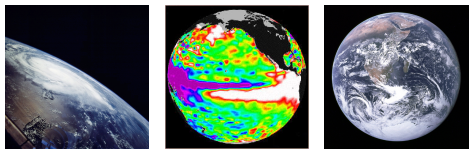

NCCS

NASA Center for Computational Sciences

Data Sharing

- Capability to share data & results
- Supports community-based development
- Facilitates data distribution and publishing

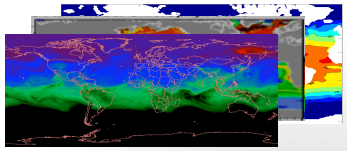

Code Development*

- Code repository for collaboration
- Environment for code development and test
- Code porting and optimization support
- Web based tools

User Services*

- Help Desk
- Account/Allocation support
- Computational science support
- User teleconferences
- Training & tutorials

DATA Storage & Management

Global file system enables data access for full range of modeling and analysis activities

Analysis & Visualization*

- Interactive analysis environment
- Software tools for image display
- Easy access to data archive
- Specialized visualization support

Data Transfer

- Internal high speed interconnects for HPC components
- High-bandwidth to NCCS for GSFC users
- Multi-gigabit network supports on-demand data transfers

HPC Compute

- Large scale HPC computing
- Comprehensive toolsets for job scheduling and monitoring

Data Archival and Stewardship

- Large capacity storage
- Tools to manage and protect data
- Data migration support

** Joint effort with SIVO*

NCCS Data Centric Climate Simulation Environment

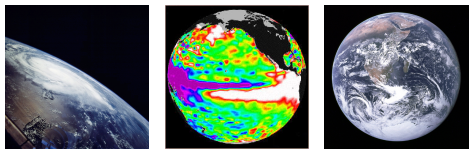

NCCS

NASA Center for Computational Sciences

Data Sharing

- Capability to share data & results
- Supports community-based development
- Facilitates data distribution and publishing

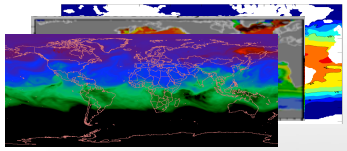

Code Development*

- Code repository for collaboration
- Environment for code development and test
- Code porting and opti… support
- Web based tools

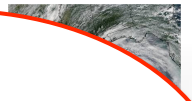

"Global" really has to be global

User Services*

- Help Desk
- Account/Allocation support
- Computational science support
- User teleconferences
- Training & tutorials

DATA Storage & Manageme…

Global file system enables data access for full range of modeling and analysis activities

- Softwa… play
- Easy access to data archive
- Specialized visualization support

Data Transfer

- Internal high speed interconnects for HPC components
- High-bandwidth to NCCS for GSFC users
- Multi-gigabit network supports on-demand data transfers

HPC Compute

- Large scale HPC computing
- Comprehensive toolsets for job scheduling and monitoring

Data Archival and Stewardship

- Large capacity storage
- Tools to manage and protect data
- Data migration support

** Joint effort with SIVO*

Notional NCCS Architecture

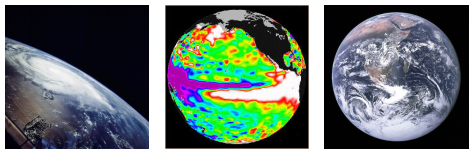

NCCS

Analysis and Visualization
Terascale *environment with tools to support interactive analytical activities*

Interactive Data Analysis

High Performance Computing
Building toward ***Petascale*** *computational resources to support advanced modeling applications*

Nehalem Cluster Upgrades

Data Storage and Management
Petabyte *online storage plus technology-independent software interfaces to provide data access to all NCCS services*

Data Archiving and Stewardship
Petabyte *mass storage facility to support project data storage, access, and distribution, access to data sets in other locations*

Data Management System

Data Portal & Earth System Grid

Data Sharing and Publication
Web-based environments to support collaboration, public access, and visualization

Notional NCCS Architecture

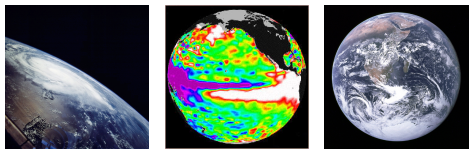

NCCS

NASA Center for Computational Sciences

Analysis and Visualization
Terascale *environment with tools to support interactive analytical activities*

Interactive Data Analysis

High Performance Computing
Building toward ***Petascale*** *computational resources to support advanced modeling applications*

Nehalem Cluster Upgrades

Data Storage and Management
Petabyte *online storage plus technology-independent software interfaces to provide data access to all NCCS services*

Data Archiving and Stewardship
Petabyte *mass storage facility to support project data storage, access, and distribution, access to data sets in other locations*

Data Management System

Internal NASA

External NASA

Data Portal & Earth System Grid

Data Sharing and Publication
Web-based environments to support collaboration, public access, and visualization

Users of Model Data

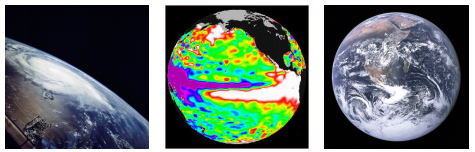

NCCS

NASA Center for Computational Sciences

Analysis and Visualization
Terascale *environment with tools to support interactive analytical activities*

High Performance Computing
Building toward ***Petascale*** *computational resources to support advanced modeling applications*

Data Storage and Management
Petabyte *online storage plus technology-independent software interfaces to provide data access to all NCCS services*

Data Archiving and Stewardship
Petabyte *mass storage facility to support project data storage, access, and distribution, access to data sets in other locations*

Data Sharing and Publication
Web-based environments to support collaboration, public access, and visualization

Non-Scientist Data Users
Access to data for national and regional policy development and decision support

Scientific Collaborators
Access to data and analysis capability by climate and weather scientific community

Integrated Access to Remote Data

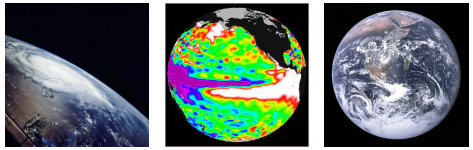

NCCS

NASA Center for Computational Sciences

Analysis and Visualization
Terascale *environment with tools to support interactive analytical activities*

High Performance Computing
Building toward ***Petascale*** *computational resources to support advanced modeling applications*

Data Storage and Management
Petabyte *online storage plus technology-independent software interfaces to provide data access to all NCCS services*

Data Archiving and Stewardship
Petabyte *mass storage facility to support project data storage, access, and distribution, access to data sets in other locations*

Remote data sources, simulation and observations

Data Sharing and Publication
Web-based environments to support collaboration, public access, and visualization

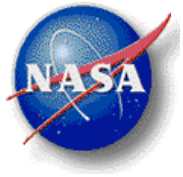

The Data Management "Problem"

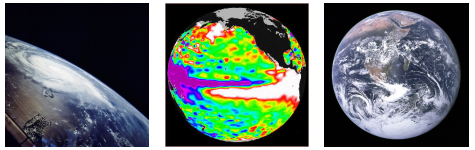

NCCS

- Genesis of the problem was the is the archive was unmanaged. No data management tools were available for users OR computing center.
- Model data was isolated in the computing center
- Costs of data was increasing each year.
  - Users could not share data
  - Users created duplicate copies
  - We backed up the duplicates
  - No method to purge old or "bad" data
  - No method to ensure that truly valuable data was protected

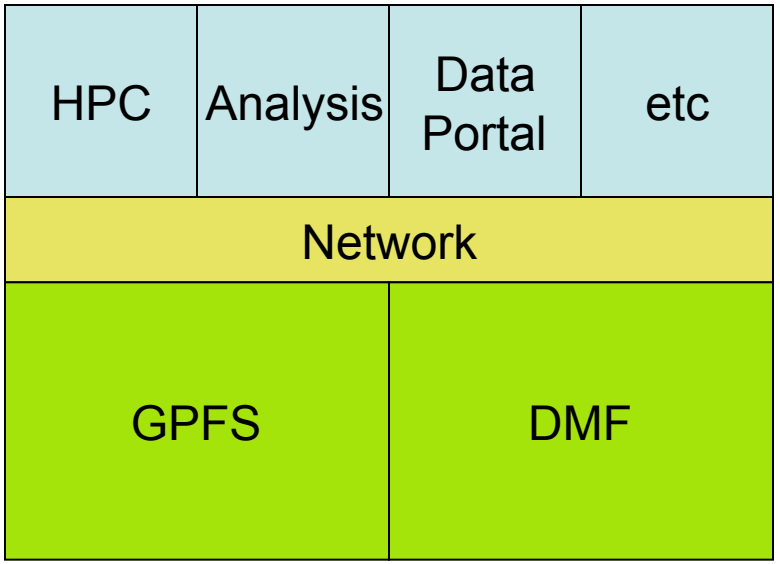

Data Management System (DMS)

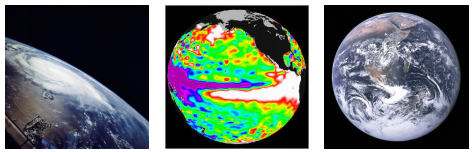

NCCS

- Developing a DMS layer in the stack.
- Based on iRODS (integrated Rule Oriented Data System) from NSF/SDSC/UNC
- Benefits
  - Ability to manage the archive and on-line storage
  - Ability to serve remote data to NCCS scientific users without importing copies
  - Ability to serve model data to other iRODS data servers.

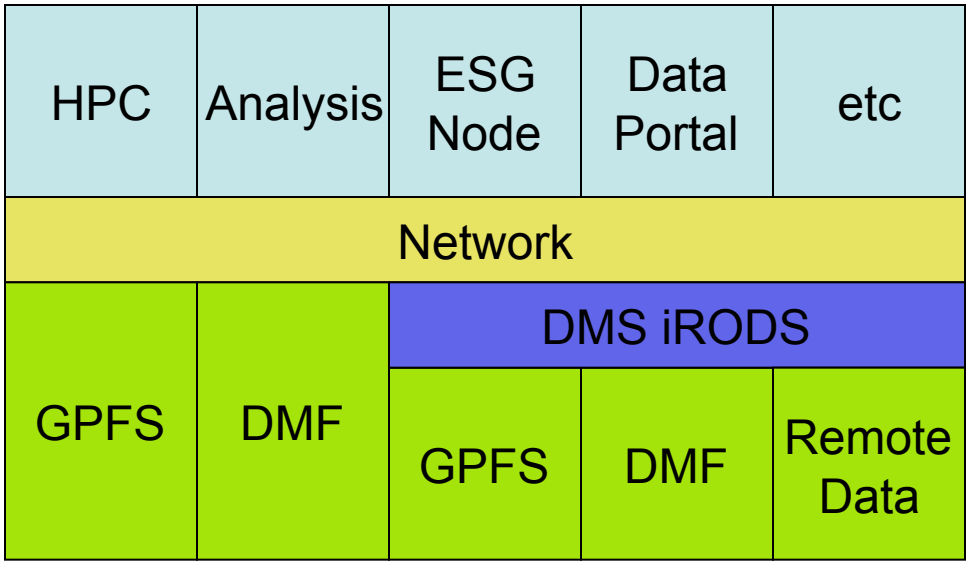

Data Management System (DMS)

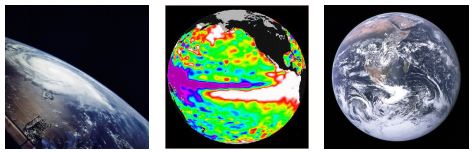

NCCS

- Developing a DMS layer in the stack.
- Based on iRODS (integrated Rule Oriented Data System) from NSF/SDSC/UNC
- Benefits
  - Ability to manage the archive and on-line storage
  - Ability to serve remote data to NCCS scientific users without importing copies
  - Ability to serve model data to other iRODS data servers.

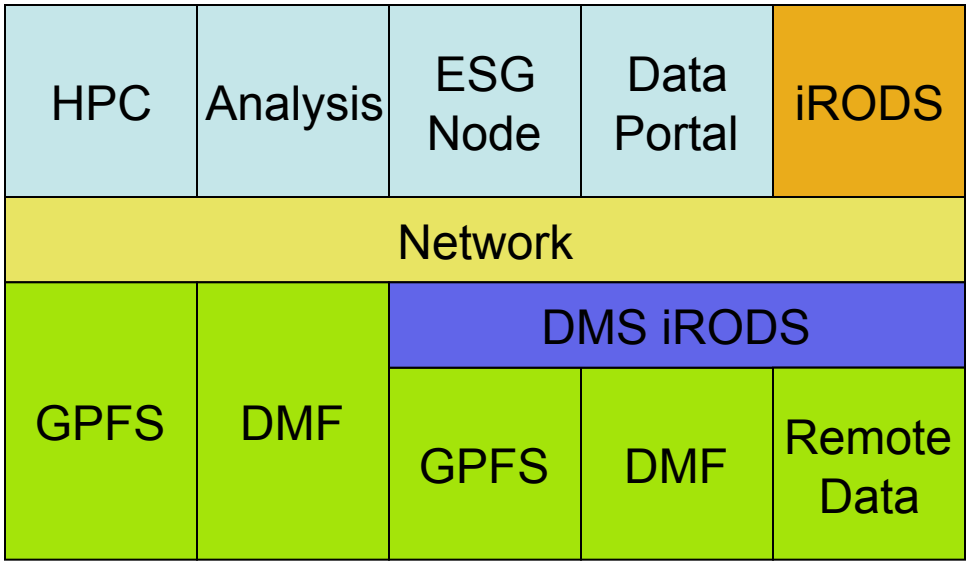

Nehalem Cluster Upgrades

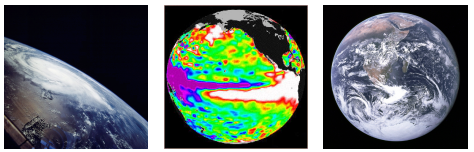

NCCS

NASA Center for Computational Sciences

- IBM iDataPlex Scalable Compute Unit (SCU) added into the Discover cluster this spring
  - 512 nodes (+46 TFLOPS)
  - 4,096 Intel Nehalem quad cores (2.8 GHz)
  - 24 GB RAM per node (+12 TB RAM)
  - Infiniband DDR interconnect
- A second 4,096 core Nehalem SCU is in Acceptance Testing
- Performance:
  - 2x speedup (per core) of some major NCCS applications
  - 3x to 4x improvement in memory to processor bandwidth
  - Dedicated I/O nodes to the GPFS file system provides much higher throughput

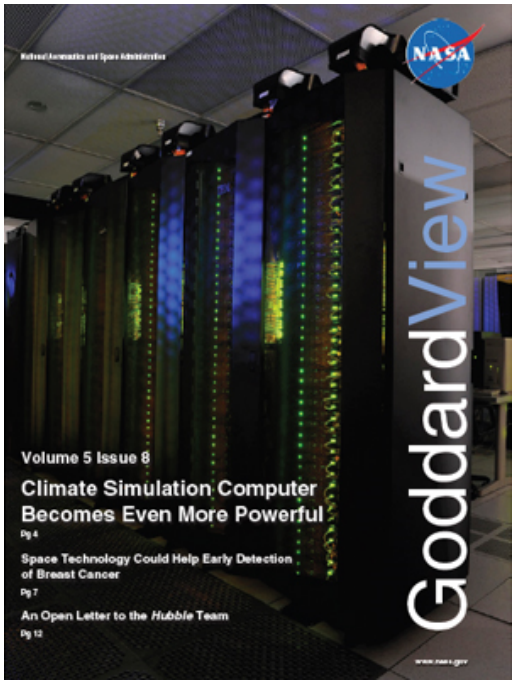

"Discover " Cluster

155 TF Peak, 14,968 cores, 34.9 TB main memory, Infiniband interconnect

• Base Unit:
- 128 nodes 3.2 GHz Xeon Dempsey (Dual Core)
• SCU1 and SCU2:
- 512 nodes 2.6 GHz Xeon Woodcrest (Dual Core)
• SCU3 and SCU4:
- 512 nodes 2.5 GHz Xeon Harpertown (Quad Core)
• SCU5 and SCU6:
- 512 nodes 2.8 GHz Xeon Nehalem (Quad Core)

Data Portal and Earth System Grid

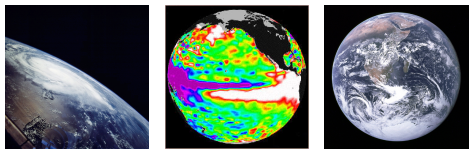

NCCS

- Web-based environments to support collaboration, public access, and visualization
- Interfaces to the Earth Systems Grid (ESG) and PCMDI for sharing IPCC model data
- Connectivity to observational data, Goddard DISC, and other scientific data sets
- Direct connection back to NCCS data storage and archive for prompt publication; minimizes data movement and multiple copies of data
- Sufficient compute capability for data analysis

NASA	Other	ESG	IPCC
Data Portal			
Local Disk	NFS	GPFS MC	iRODS

Data Portal Platform
(128 cores, 1.2TF, 120TB of disk)

Earth System Grid

Interactive Data Analysis & Visualization Platform

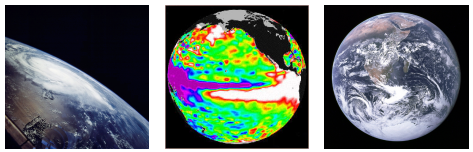

NCCS

NASA Center for Computational Sciences

- Interactive Data Analysis Systems
  - Direct login for users
  - Fast access to all file systems
  - Supports custom and 3rd party applications
  - Visibility and easy access to post data to the data portal
  - Interactive display of analysis results
- In-line and Interactive visualization
  - Synchronize analysis with model execution
  - Access to intermediate data as they are being generated
  - Generate images for display back to the user's workstations
  - Capture and store images during execution for later analysis
- Develop Client/Server Capabilities
  - Extend analytic functions to the user's workstations
  - Data reduction (subsetting, field/variable/temporal extractions, averaging, etc.) and manipulation (time series, display, etc.) functions

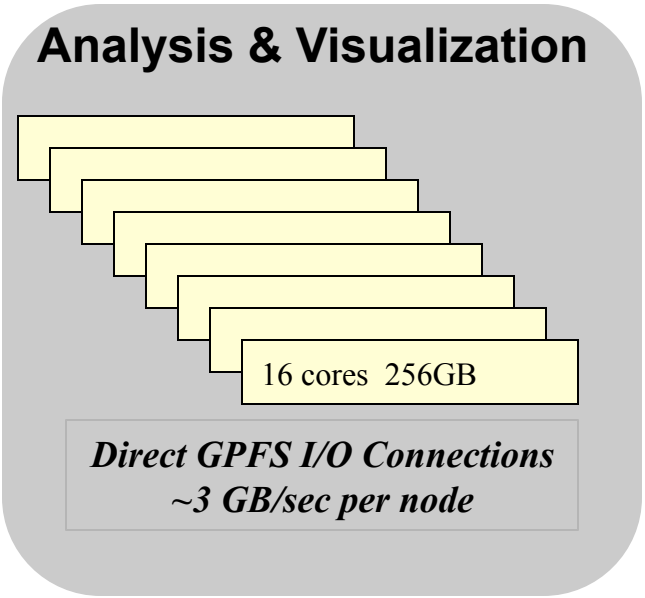

Dali Analytics Platform

1.2 TF Peak, 128 cores, 2 TB main memory

- 8 nodes 2.4 GHz Dunnington (Quad Core)
- 16 cores/node with 256 GB memory/core
- 3 GB/s I/O bandwidth to GPFS filesystem
- Software: CDAT, ParaView, GrADS, Matlab, IDL, python, FORTRAN, C, Quads, LATS4D

Currently configured as (8) 16-core nodes with 256 GB RAM/node, with flexibility technology to support up to (2) 64-core nodes with 1 TB RAM/ node.

Conclusion

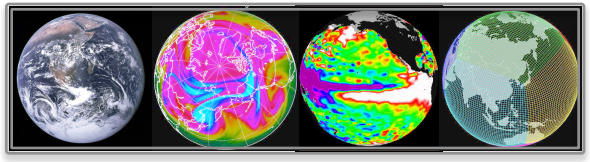

- NASA modeling efforts are advancing in resolution and complexity for both weather and climate prediction, realizing
  - Improved representation of nature
  - Improved utilization of NASA satellite data
- Increased compute capacity enables higher resolution even for the ensembles needed to characterize uncertainty in climate prediction and projection
- The NCCS is moving forward to support NASA's climate & weather research that will
  - Enhance scientific value of current satellite observations
  - Accelerate readiness for upcoming Decadal Survey* missions
  - Support international collaborative projects and national applications

*January 2007 NRC report: *Earth Science and Applications from Space: National Imperatives for the Next Decade and Beyond.*